UAT-WP-19-05
12 January 2004

<u>RTCA Special Committee 186, Working Group 5</u>

<u>ADS-B UAT MOPS (DO-282), Revision A</u>

<u>Meeting #19</u>

<u>Teleconference on 1.12.04</u>

Proposal to add an Appendix [P], entitled "UAT Message Overlap Statistics" to the UAT MOPS

Prepared by Warren J. Wilson and Larry Bachman

SUMMARY

This Working Paper addresses Action Item 17-02, which called for the review of Working Paper UAT-WP-10-01 authored by Warren Wilson. Issues related to the number of overlapping ADS-B signals that need to be accommodated are discussed. Several approaches to evaluating the number and frequency of overlapping message statistics are presented, along with the resulting analysis. The main finding is that a receiver that can handle three simultaneous overlapping signals will be able to receive the vast majority of cases.

1. Introduction

This Appendix addresses issues pertaining to the statistics of overlapping ADS-B signals in a multi-user environment. These issues are of interest because they have an impact on receiver design requirements and, hence, on MOPS test procedures. The basic question is to determine how many overlapping ADS-B signals a UAT receiver must be able to process. This question is best answered in the context of a hypothetical receiver architecture. The receiver design considered here is shown in Figure P-1.

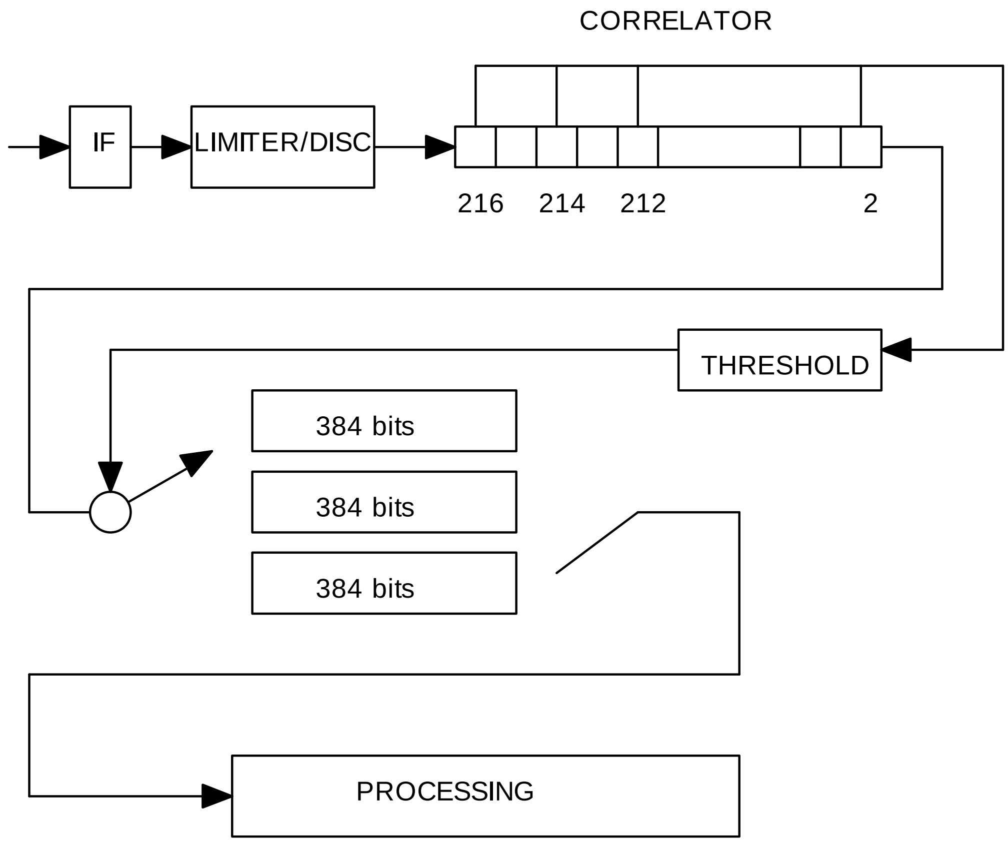

Figure P-1: Possible UAT Receiver Architecture

In this design the incoming signal is demodulated into a string of ones and zeroes that is compared to a known synchronization sequence in a correlator. A sequence that passes a certain threshold will cause the ensuing bit samples to be placed in one of a number of 384-bit registers. The size of the registers is determined by the size of a long ADS-B message. Such messages are sent as RS(48, 34) code words, so the required length of a register is just 48 x 8 = 384 bits. It is assumed that when one of these registers is filled its contents are immediately passed along to the input queue of a RS decoder and the register is available for another incoming message. Of course, it is possible that the incoming message is a basic ADS-B message whose length is only 240 bits (since it is a RS(30, 18) code word). The system design is based on the ability of the RS decoder to sort out (with very high probability) which of the two possible types was actually sent. It is also possible that the RS decoder will determine that the message cannot be decoded because there are too many errors or because the whole message was the result of a false alarm.

The issue to be resolved in this appendix can be paraphrased as, "How many registers are necessary?" The need for more than one register arises whenever a synchronization occurs while a register is still being filled with bits derived from a previous synchronization. Such a situation could happen for a number of reasons:

1. A new signal which is significantly stronger than the first signal arrives.
2. The signal being demodulated contains a sequence of bits that is sufficiently like the synchronization sequence to pass the threshold (embedded sync).
3. The first signal was a basic ADS-B message, and a new message arrives soon after it is completed.

The receiver has no way of knowing which of these conditions applies in any given situation. If the receiver *did* know it might be able to take some appropriate action, which would obviate the need for more than one register. For example, if it were known that case 1 applied, the best course of action would be to replace the old message with the new one. The second case could occur if some of the information in the message "looked like" the synchronization sequence. This similarity might be relatively static (for example, a way point might persist for long periods of time). Thus, if this case applied, the initial message should not be removed. In the third case, the receiver could pass along an abbreviated message with instructions to the RS decoder to attempt only RS(30, 18) decoding. Because the appropriate actions are much different in each case, it is highly desirable to avoid making a choice prior to the RS decoding process.

2. Statistical Model Results

This section will describe a theoretical statistical model used to derive an upper bound for the number of registers required to provide a given probability of message processing. Section 3 will present the results of using the Multi-Aircraft UAT Simulation (MAUS) to evaluate the message overlaps in the high-density Core Europe 2015 air traffic environment.

A set of scenarios was modeled in which a number (N) of users were assumed to be transmitting ADS-B messages. Although it is expected that some fraction of the messages will actually be basic ADS-B messages, a full complement of long messages was assumed in order to provide a worst-case analysis. From the point of view of any particular receiver, the arrival times of the messages was assumed to be uniformly distributed over the 800 ms portion of each second devoted to ADS-B. The power levels of the messages were chosen randomly from the reverse cumulative distribution shown in Figure P-2. This distribution was derived from MAUS results during the spring of 2001 and was based on the UAT deployment assumptions being used at the time. The markers represent the simulation data, and the line is based on the equation that was used in the model:

$$P(x) = \exp(-0.0046\,(x + 104.7)^2) \qquad (1)$$

for $x \geq -104.7$. $P(x) = 1$ for $x < -104.7$. In this equation, x is measured in units of dBm. The results of this analysis are *not* very sensitive to the exact form of the distribution.

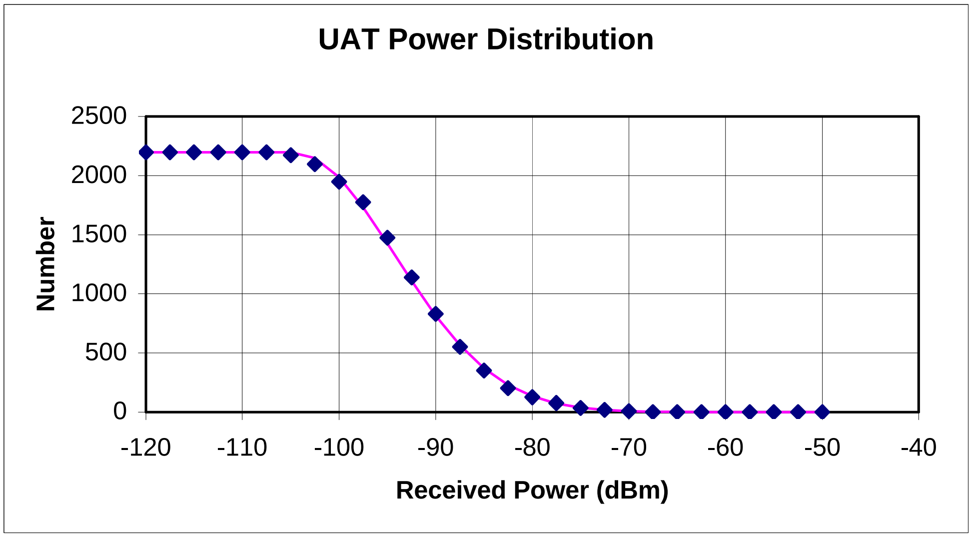

Figure P-2: UAT Received Power Distribution (for 2200 users)

A very simplified version of the model assumes that the synchronization portion of each message always succeeds and every message needs to be placed in one register or another. The register requirements would then depend on the arrival sequence of the messages. A number of possibilities are shown in Figure P-3.

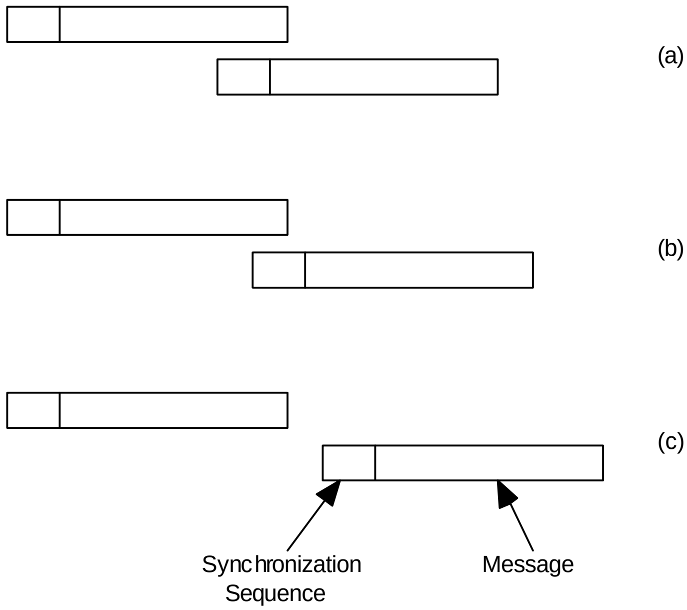

Figure P-3: Message Overlap Possibilities

In case (a) of Figure P-3, the message portions of two bursts are coincident in time, and they require separate registers. If an additional register is not available, the second message is dropped. In case (b), the message portion of one burst overlaps the synchronization portion of another; thus, they can each use the same register. In case (c), the bursts are entirely separate, and one register will clearly suffice.

In this simplified case the probability of requiring a certain number of registers can be calculated by noting that one message can take precedence over another if it precedes it by a number of bit periods between 1 and 384. If there are N users competing with a given burst, the average number of preceding users within the last 384 bit periods is given by

$$l = 384 \times N / 833334 \qquad (2)$$

where 833334 is the number of bit periods in 0.8 seconds. The probability of requiring n registers is given by a Poisson distribution, i.e.,

$$P(n) = e^{-l}\, l^{n-1} / (n-1)! \qquad n \geq 1 . \qquad (3)$$

In the real world the probabilities (as shown by the MAUS results presented in Section 3) will be altered because many of the overlapped signals will not synchronize. For example, many times the second signals in Figure P-3(a) and Figure P-3(b) will not synchronize because of interference from the first signal. It is necessary to have a model of synchronization performance to account for these effects. The simplified model of synchronization performance used in this theoretical analysis is as follows:

1. If the desired signal power is greater than the power of any individual interferer overlapping the synchronization sequence, then the synchronization will succeed.
2. If the desired signal power is less than the power of any interferer overlapping its synchronization sequence, then the synchronization will fail if and only if at least 16 bits of the synchronization sequence are overlapped. [The number 16 relates to a synchronization threshold value of 84 (out of a possible 108). When there are 16 bits overlapped by a strong signal (producing a 50% BER), the average synchronization score is 3x20 + 3x16/2 = 84.]

These criteria overestimate the probability of synchronization. The criteria fail to properly take into account cases where there are multiple interferers. Also, criterion 1 is clearly optimistic even when there is only one interferer. Overestimating the synchronization probability will once again tend to overbound the number of registers required.

With these rules in place, the number of registers required to accommodate all potential messages can be determined by simulation. The results for various values of N are plotted in Figure P-4 through Figure P-11 as the curves labeled “real.” These show the average number of messages per second that require the specified number of registers. For comparison, the results of equation (3) are shown as the curves labeled “ideal.”

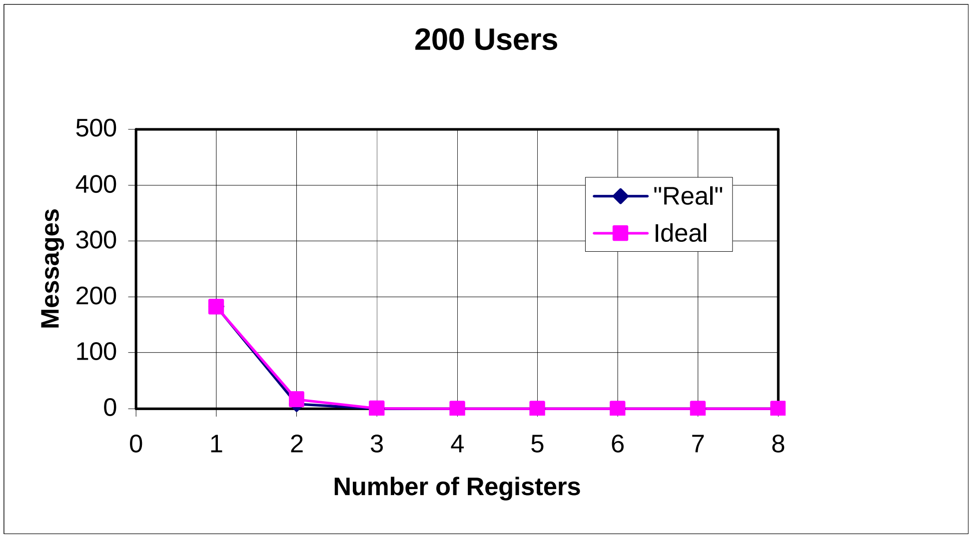

Figure P-4: Simulation Results for 200 Aircraft

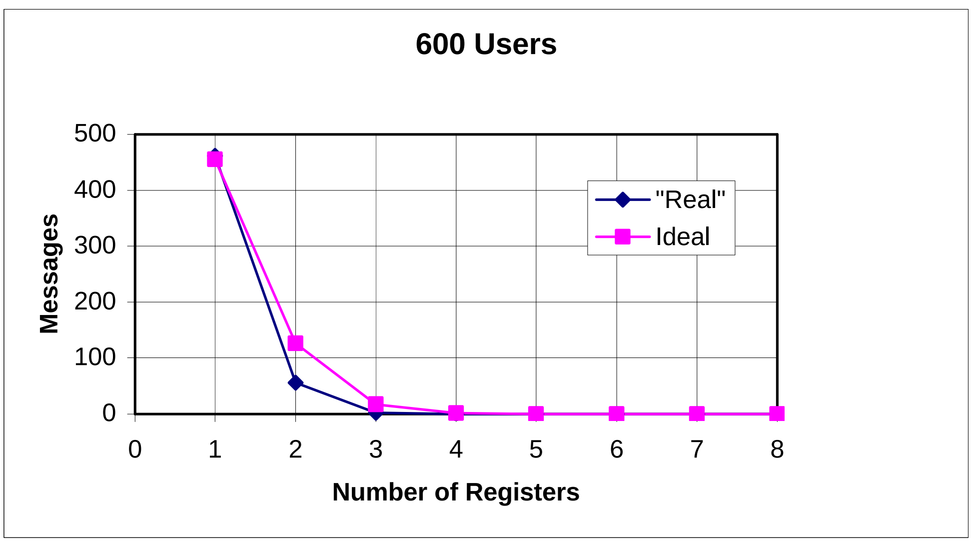

Figure P-5: Simulation Results for 600 Aircraft

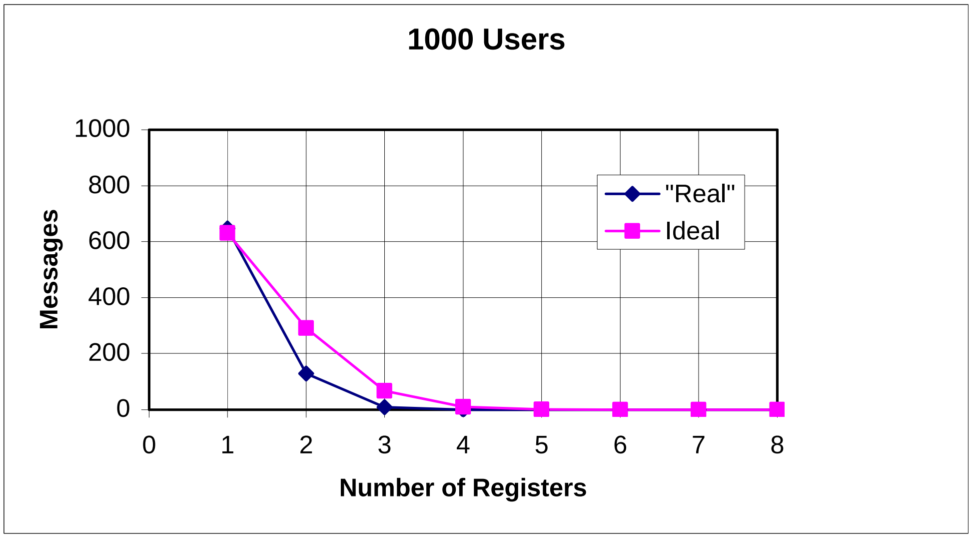

Figure P-6: Simulation Results for 1000 Aircraft

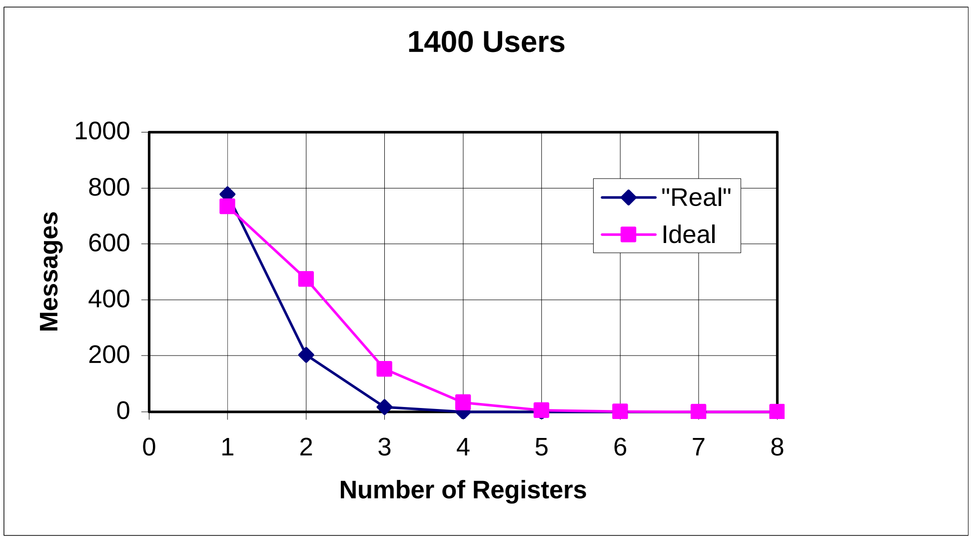

Figure P-7: Simulation Results for 1400 Aircraft

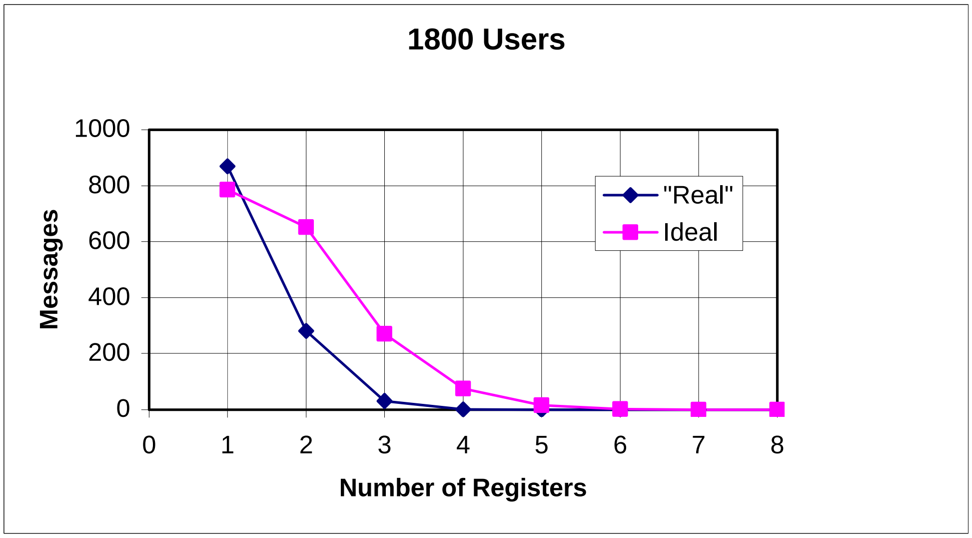

Figure P-8: Simulation Results for 1800 Aircraft

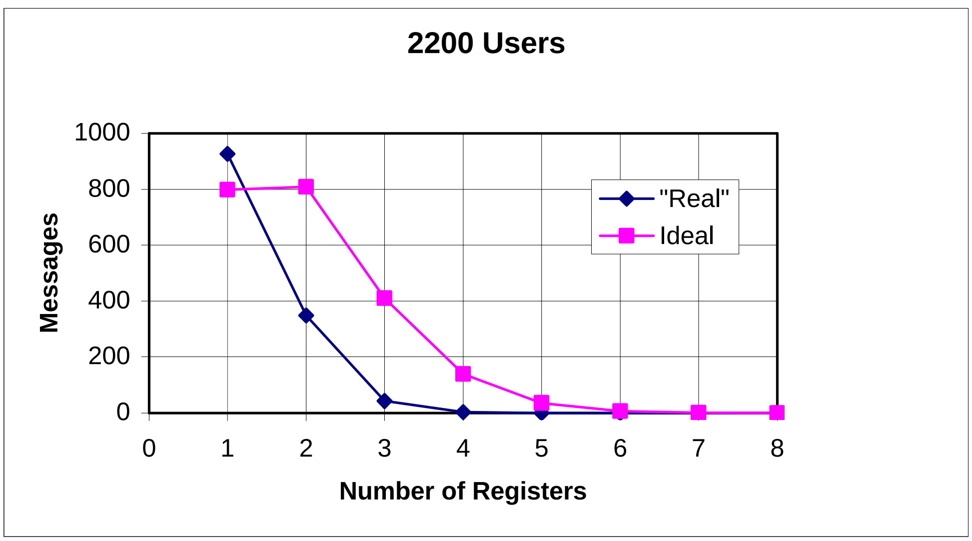

Figure P-9: Simulation Results for 2200 Aircraft

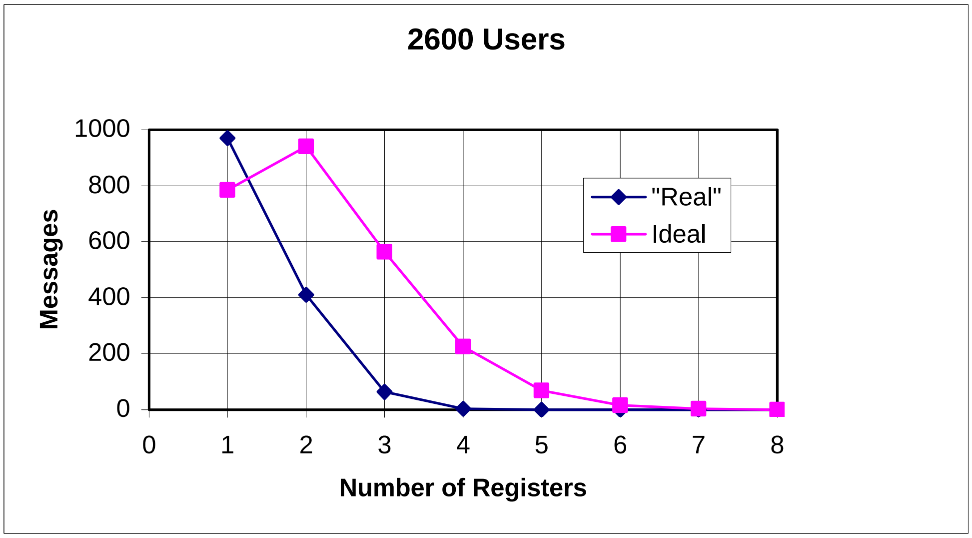

Figure P-10: Simulation Results for 2600 Aircraft

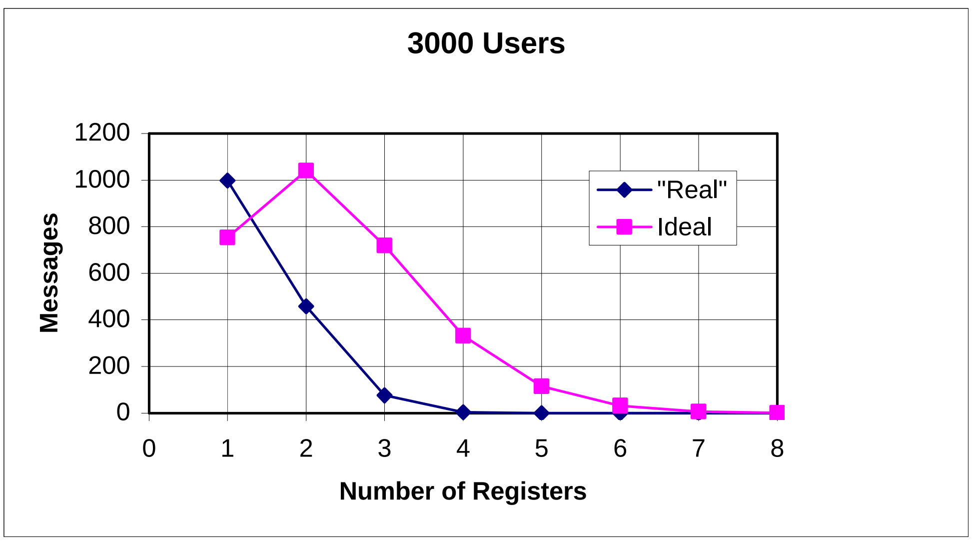

Figure P-11: Simulation Results for 3000 Aircraft

2.1. Discussion

There are various trends that can be seen in the graphs. For instance, as the number (N) of users increases, the fraction of the messages that actually synchronize decreases. The number of successful synchronizations versus N is plotted in Figure P-12. This is evidence that the system is beginning to saturate. These graphs do not, however, take into account the messages that successfully synchronize but do not survive the error correction process. That information is irrelevant to this discussion because it has no impact on the number of registers used.

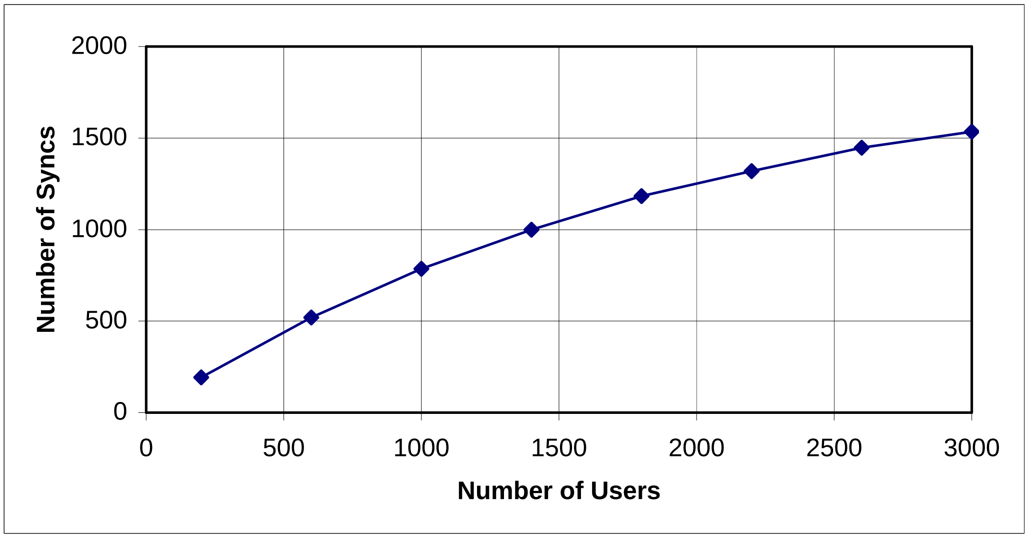

Figure P-12: Overall Synchronization Performance

A second trend that can be seen in the results is that the "real" distributions tend to be much more peaked at the value 1 than the ideal curves. This happens because the signals with multiple overlaps are the ones most likely to be pruned due to synchronization failures. More realistic models of synchronization performance would probably result in distributions with even more peaking. It is interesting to note that as the number of users increases the shapes of the "real" distributions become very similar.

The most important issue for this section is the fraction of potentially successful messages that are rejected if the number of registers is limited. In Figure P-13, this fraction is plotted versus N, assuming 1, 2, or 3 registers are available. The curves tend to flatten out as the number of users increases. This seems to be due to the fact that the shapes of the curves become similar as N increases. (The irregularities in the curves are related to the limited number of cases studied. These curves are based on simulations of 10 seconds' worth of UAT transmissions.) For N = 2200, which is the number used for the LA 2020 scenario, the values are 0.30 for one register, 0.034 for two registers, and 0.0022 for three registers. The use of three registers reduces the probability of rejection to a negligible level.

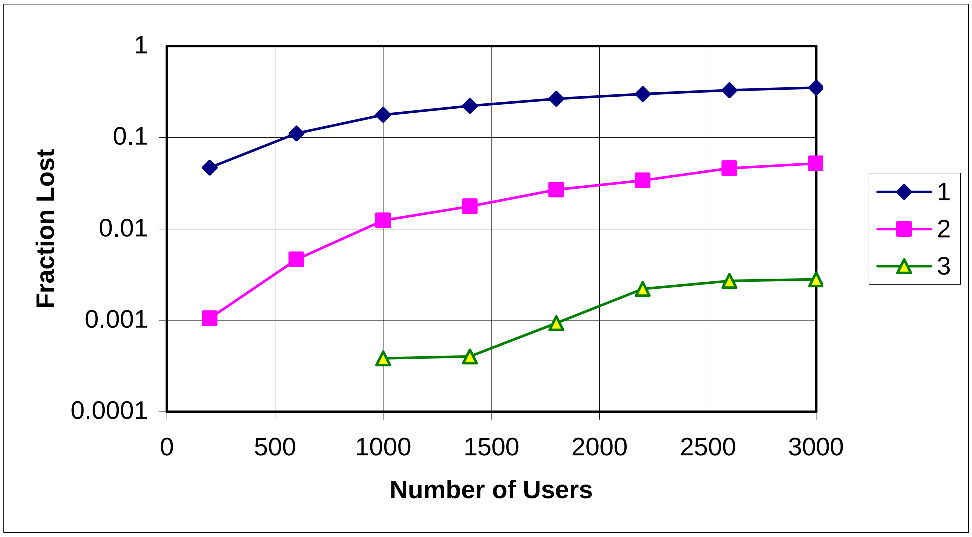

Figure P-13: Fraction of Rejected Messages (Parameter is the Number of Registers)

2.2. Effect of Uplink Messages

In the normal operation of the UAT system, the ADS-B bursts and the uplink bursts are supposed to be transmitted in different time intervals separated by relatively large guard times. Nevertheless, there is a system requirement to continually search for both types of messages simultaneously. If that is the case, then there is some possibility that a receiver can "synchronize' with an uplink message during the time usually reserved for ADS-B messages. This could happen due a false alarm or because something close to the uplink synchronization sequence is embedded in an actual ADS-B message. Either way, the potential uplink message will occupy some receiver resources that would otherwise be available for ADS-B.

Suppose, for example, that at least one of the registers shown in Figure P-1 is actually 4416 bits in length instead of 384. Such a register is just long enough to contain an entire up link message (4416 = 6 x 8 x 92). Whenever there is an uplink synchronization detected the message is placed in such a register, which is then not available for duty receiving ADS-B information. Thus, if the total number of registers is 3, then the effective number available for ADS-B is reduced to 2 for as long as it takes to fill the long memory. If the uplink threshold is set at a value that yields about 1 false alarm or embedded synchronization per second, then the average fraction of time spent on extraneous uplink message processing is 4416/833334 = 0.0053. This can be converted into a revised estimate of the probability of rejecting an ADS-B message. For N=2200 and a total of three registers, the revised probability is 0.0053 x 0.034 + 0.9947 x 0.0022 = 0.0024. In other words, the effects of false uplink synchronization phenomena appear to be negligible if the false alarm rate is reasonably small.

3. MAUS Results

The MAUS (see Appendix K) was run in the Core Europe 2015 air traffic environment to determine the distribution of message overlaps on ADS-B messages being received. MAUS looks at the arriving message stream at the receiver and combines the incoming ADS-B messages, bit by bit, with all other sources of signals (e.g., other ADS-B messages, interference from Link 16 and DME, on-board transmissions) to feed into the receiver performance model. The receiver performance model then computes a bit error rate for each bit and applies the appropriate Reed-Solomon error correction to the bits of the message to determine a message error rate (MER). In the process of calculating this MER, the model must determine the probability of reception of the synchronization sequence. It is this probability that was used to examine the distribution of message overlaps.

Figure P-14 shows a distribution of probability of error in receiving the synchronization sequence of an incoming ADS-B message for an A2 equipage receiver at FL 400 in the Core Europe 2015 scenario. Note the peaks at 0 and 1, with 1 representing messages with no chance of synch reception due to strong interference and 0 symbolizing messages that arrive with little or no simultaneous interference present. Values between 0 and 1 may correspond to messages that arrive with only partial interference overlap during the synchronization sequence, for example. Since the majority of cases fall at the extremes of 0 or 1, it was decided to place the cutoff between success and failure at a probability of successful synch reception of 0.5. The results are not sensitive to this choice.

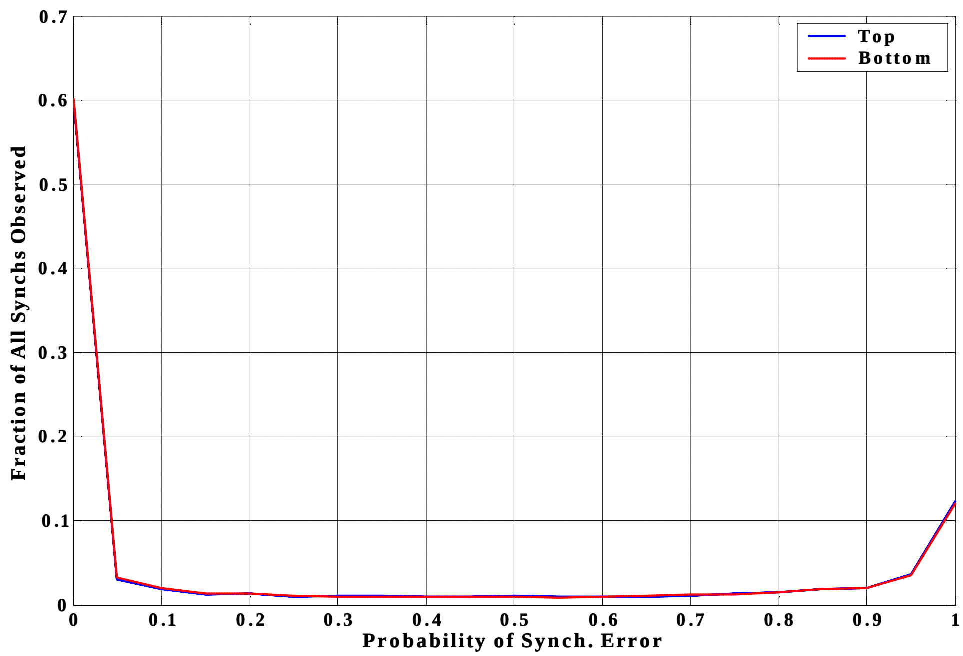

Figure P-14: Probability of Error in Reception of the Synchronization Sequence in Core Europe 2015

For the purposes of this appendix, the analysis focused on the successful synch receptions. The metric that was used was the distribution of the number of successful synchs that occur while the receiver is attempting to decode a previously synched message. This metric was chosen, since the assumption was made that only subsequent successful synch detections during message reception would require additional registers. This assumption ignores unsuccessful synch arrivals, since no register would be required if the synch was not successfully received. (This is true even though the unsuccessful message could interfere with reception of the message associated with the previous successfully received synch.) Therefore, the metric is simply a count of successful synch detections, which arrive within the next 420 usec after an initial successful synch detection (for the long ADS-B message).

Figure 15 shows the distribution of the reception of successful synchs during message reception following a previous successful synch. The overlap distribution is shown for an A2 equipped aircraft in the Core Europe 2015 air traffic scenario, flying at FL 400. These results indicate that more than 99.9% of the receivable messages have two or fewer overlapping ADS-B messages.

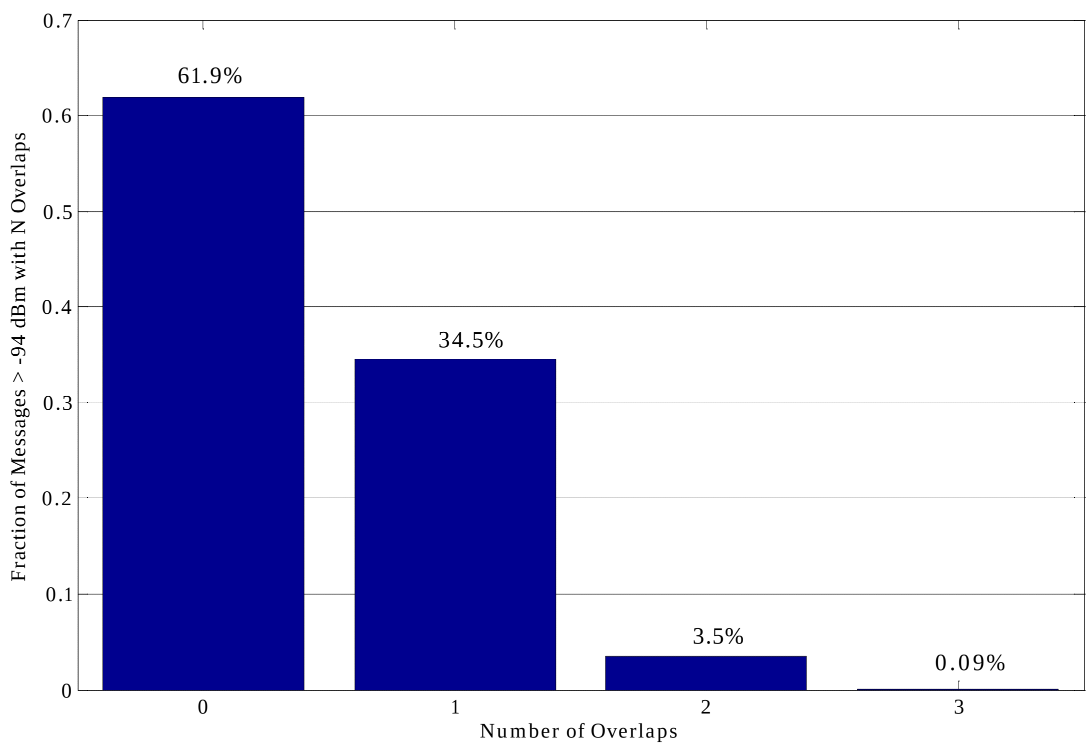

Figure P-15a: Percentage of Successful Overlapping Synch Sequences for Messages Received > -94 dBm for an A2 at 40,000 ft. in CE 2015 for the Bottom Antenna

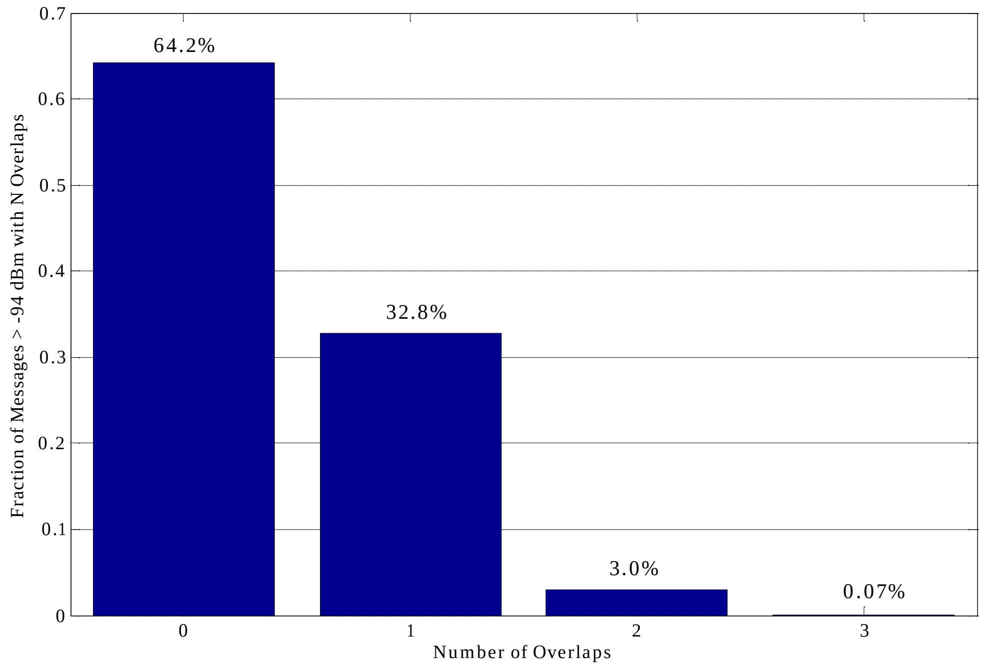

Figure P-15b: Percentage of Successful Overlapping Synch Sequences for Messages Received > -94 dBm for an A2 at 40,000 ft. in CE 2015 for the Top Antenna

The results shown in Figure P-15, while differing somewhat from those determined by the theoretical analysis of Section 2 (which had many simplifying assumptions), confirm the conclusion reached in that section, namely that a receiver that uses three registers will be capable of decoding the vast majority of ADS-B messages.

4. Summary

Performance has been simulated using both a simplified system performance model based on using a separate memory register to process each incoming message, and by the Multi-Aircraft UAT Simulation, which performs a more detailed simulation of the high-density environment and is based on measured receiver performance. When the number of messages being received exceeds the number of registers, new messages are dropped. Using this rule, it appears that the ability to handle three overlapped signals will cover almost all cases. (This is not meant to imply that an actual UAT implementation needs to be designed this way.)